arXiv:1703.06359v1 [math.NA] 18 Mar 2017

FULLY SYMMETRIC KERNEL QUADRATURE*

TONI KARVONEN† AND SIMO SÄRKKÄ†

Abstract. Kernel quadratures and other kernel-based approximation methods typically suffer from prohibitive cubic time and quadratic space complexity in the number of function evaluations. The problem arises because a system of linear equations needs to be solved. In this article we show that the weights of a kernel quadrature rule can be computed efficiently and exactly for up to tens of millions of nodes if the kernel, integration domain, and measure are fully symmetric and the node set is a union of fully symmetric sets. This is based on the observations that in such a setting there are only as many distinct weights as there are fully symmetric sets and that these weights can be solved from a linear system of equations constructed out of row sums of certain submatrices of the full kernel matrix. We present several numerical examples that show feasibility, both for a large number of nodes and in high dimensions, of the developed fully symmetric kernel quadrature rules. Most prominent of the fully symmetric kernel quadrature rules we propose are those that use sparse grids.

Key words. Numerical integration, kernel quadrature, Bayesian quadrature, reproducing kernel Hilbert spaces, fully symmetric sets, sparse grids

AMS subject classifications. 46E22, 47B32, 60G15, 65C05, 65C50, 65D30, 65D32

1. Introduction.

Let Ω be a subset of $\mathbb{R}^d$, μ a measure on Ω, and $f\colon \Omega \to \mathbb{R}$ a function that is integrable with respect to μ. Computation of the integral $\mu(f) \coloneqq \int_\Omega f \,\mathrm{d}\mu$ is a recurring problem in applied mathematics and statistics. In most cases this integral has no readily available analytical form and one must resort to a *quadrature rule* (or, occasionally, a *cubature rule* if $d > 1$) for its approximation. A quadrature rule Q is a linear functional of the form

$$Q(f) \coloneqq \sum_{i=1}^{n} w_i f(\mathbf{x}_i) \approx \int_\Omega f \,\mathrm{d}\mu,$$

where $\mathbf{x}_i \in \mathbb{R}^d$ are the *nodes* and $w_i \in \mathbb{R}$ are the *weights* that greatly affect accuracy of the quadrature rule.

The nodes and weights are usually chosen so that the quadrature approximation is exact whenever the integrand is a low-degree polynomial [10, 9]—such methods are called *classical* or *polynomial* quadrature rules in this article. Another possibility is to use Monte Carlo or quasi quasi Monte Carlo methods [6].

Here we study *kernel quadrature rules* that are, for arbitrary fixed nodes, optimal (in the sense of minimising the worst-case error) quadrature rules in reproducing kernel Hilbert spaces (RKHS). Theory of these rules goes back at least to the work of Larkin [28, 29] in the 1970s. There is also an average-case interpretation of these quadrature rules [51]. Lately, kernel quadrature rules have been a subject of renewed interest because they can be used for numerical integration on scattered data sets [2, 53] and they carry a probabilistic interpretation as posterior means for Gaussian processes assigned to the integrand [43]. The probabilistic interpretation, equivalent to the RKHS formulation we use, is interesting because it allows for statistical modelling of error in numerical integration and is one of main motivators behind this article. The above topics, including the probabilistic interpretation, are reviewed in Section 2.

*Submitted to arXiv on March 18, 2017.

Funding: This work was supported by Aalto ELEC Doctoral School as well as Academy of Finland projects 266940 and 273475.

†Department of Electrical Engineering and Automation, Aalto University, Espoo, Finland (toni.karvonen@aalto.fi, simo.sarkka@aalto.fi).

An advantage of kernel quadrature rules is that the nodes are *not* prescribed as opposed to polynomial quadrature rules (polynomial rules with arbitrary nodes could probably be developed along the lines of de Boor and Ron interpolation [11, 35]). The flexibility comes with the price of having to solve the weights from a linear system of n equations, a task of cubic time and quadratic space complexity. It would not be practical to tabulate the weights beforehand as they depend on the kernel, integration domain, and measure. Partially due to the computational cost, only low numbers of nodes have been used in kernel quadrature and much of the literature is concerned with optimal selection of the nodes. See [28, 43, 44, 33, 54] for some optimal node configurations and [3] for an approximate alternative for finding "good" nodes. Efficient computation of the optimal nodes is an open problem and not the topic of this article. Instead, we select the nodes such that the weights can be easily computed.

There is little work for extending applicability of kernel quadrature to integration problems where it is necessary to use a large number of nodes due to high dimensionality or high level of accuracy that is desired. O'Hagan [43, 44] has proposed some computationally beneficial node designs that are unfortunately too inflexible to be of any use in most situations. Many fast and approximate kernel-based methods have been developed in scattered data approximation, statistics, and machine learning literature (see e.g. [4, Supplement C] for a compendium). However, accuracy of quadrature rules is often strongly dependent on the weights having been computed exactly. Furthermore, approximate weights give rise to some philosophical objections if they are to be used also for statistical modelling of error in an integral approximation.

In this article we show that if certain structure is imposed on the node set, then the kernel quadrature weights can be computed *exactly* and in a very simple manner. Our approach is based on *fully symmetric sets* [16, 17] which are point sets that can be obtained from a given vector through permutations and sign changes of its coordinates. In Section 3 we show that some symmetry assumptions on Ω and μ (see Assumption 3.4) lead to, for a certain large class of kernels that includes all isotropic kernels, tractable computation of the weights if the node set is a union of fully symmetric sets. Depending on the dimension, the weights can be computed for sets of this type that contain of up to tens of millions of nodes. The crucial observation under our assumptions is that *there are only as many distinct weights as there are fully symmetric sets making up the node set*. The *fully symmetric kernel quadratures* we construct exhibit the following advantageous properties:

- The algorithm (part of Theorem 3.6) for exact computation of the weights is exceedingly simple and easy to implement.
- If there are J fully symmetric sets containing n nodes in total, only Jn kernel evaluations are needed. In all realistic situations J is at most a few hundred while n can, as mentioned, go up to millions (Section 5.4 has an example where $J = 832$ and $n = 15{,}005{,}761$). The weights are solved from a system of J linear equations and only a $J \times J$ matrix needs to be stored.
- The node selection scheme remains quite flexible and the number of nodes does not grow too fast with the dimensions (see Equation (3.2) and Table 4.1) as happens when for example full Cartesian grids are used. The smallest non-trivial fully symmetric sets contain $2d$ points.

In Section 4 we discuss a number of possible ways of selecting the fully symmetric sets. Sparse grids [5], popular in polynomial-based high-dimensional quadrature, seem the most promising alternative. We also provide some theoretical convergence guarantees (Theorem 4.2) for kernel quadratures that use Clenshaw–Curtis sparse grids [37]. Some numerical experiments are carried out in Section 5.

The fast weight algorithm for computing the weights is not easily extended to fitting of the kernel parameters that often have considerable effect on accuracy of the integral approximation. Our experiments show that fully symmetric kernel quadratures are feasible but we have to resort to *ad hoc* solutions for fitting the kernel parameters or know them beforehand. Development of efficient fitting procedures is left for future research. This is discussed in Section 5.1.

Finally, we remark that this is not the first instance of fully symmetric sets being used in conjunction with kernel quadrature. In the context of Gaussian filtering of non-linear systems, kernel quadrature version with the simplest non-trivial fully symmetric set (called in filtering literature the unscented transform) has seen use in [54, 46, 47]. These articles do not use the weight computation algorithm that we present.

2. Kernel quadrature.

This section reviews the basics of quadrature rules in reproducing kernel Hilbert spaces. We also briefly discuss connections to probabilistic modelling of numerical algorithms. See for example [28, 4] for proofs and additional references. Standard reference for reproducing kernel Hilbert spaces is [1].

2.1. Quadrature in reproducing kernel Hilbert space.

Every positive-definite kernel $k\colon \Omega \times \Omega \to \mathbb{R}$ defines a *reproducing kernel Hilbert space* $\mathcal{H}$ of functions $f\colon \Omega \to \mathbb{R}$ through the properties (i) $k(\cdot, \mathbf{x}) \in \mathcal{H}$ for every $\mathbf{x} \in \Omega$ and (ii) pointwise evaluations of any $f \in \mathcal{H}$ can be represented in terms of inner product with the kernel: $\langle f, k(\cdot, \mathbf{x})\rangle_{\mathcal{H}} = f(\mathbf{x})$. The integral operator μ and the quadrature rule Q are bounded linear functionals on $\mathcal{H}$ under the non-restrictive assumption $\int_\Omega \sqrt{k(\mathbf{x}, \mathbf{x})}\,\mathrm{d}\mu(\mathbf{x}) < \infty$. The *worst-case error* (WCE) $e(Q)$ of a quadrature rule Q is defined in terms of the dual norm

$$e(Q) \coloneqq \|\mu - Q\|_{\mathcal{H}^*} = \sup_{\|f\|_{\mathcal{H}} \leq 1} |\mu(f) - Q(f)| \tag{2.1}$$

that can be also written as $e(Q) = \|\mu[k(\cdot, \mathbf{x})] - Q[k(\cdot, \mathbf{x})]\|_{\mathcal{H}}$. Why this is a reasonable measure of error of the quadrature rule is apparent after an application of the reproducing property and the Cauchy–Schwarz inequality:

$$|\mu(f) - Q(f)| = \big|\langle f, \mu[k(\cdot, \mathbf{x})] - Q[k(\cdot, \mathbf{x})]\rangle_{\mathcal{H}}\big| \leq e(Q)\,\|f\|_{\mathcal{H}}$$

for $f \in \mathcal{H}$. That is, if the integrand belongs to the RKHS, convergence in the usual sense of diminishing integration error is implied by convergence to zero of the WCE. Relationship between the kernel and its induced RKHS is further discussed in Section 4.4 where we also provide two convergence theorems for the worst-case error.

The quadrature rule that, for arbitrary fixed distinct nodes $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$, minimises the worst-case error (2.1) is called the *kernel quadrature rule* and denoted by Q_k. This rule is unique and the optimal weights $\mathbf{w} = (w_1, \ldots, w_n)$ can be solved from

$$\underbrace{\begin{pmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & \cdots & k(\mathbf{x}_1, \mathbf{x}_n) \\ \vdots & \ddots & \vdots \\ k(\mathbf{x}_n, \mathbf{x}_1) & \cdots & k(\mathbf{x}_n, \mathbf{x}_n) \end{pmatrix}}_{=\mathbf{K}} \begin{pmatrix} w_1 \\ \vdots \\ w_n \end{pmatrix} = \underbrace{\begin{pmatrix} \mu_k(\mathbf{x}_1) \\ \vdots \\ \mu_k(\mathbf{x}_n) \end{pmatrix}}_{=\mu_k(\mathcal{X})}, \tag{2.2}$$

where $\mathbf{K}$ is the positive-definite *kernel matrix* and $\mu_k(\mathbf{x}) \coloneqq \int_\Omega k(\mathbf{x}, \mathbf{x}')\,\mathrm{d}\mu(\mathbf{x}')$ the *kernel mean*, an object of much independent interest [34]. The kernel quadrature rule

and its worst-case error are

$$Q_k(f) = \sum_{i=1}^{n} [\mathbf{K}^{-1}\mu_k(\mathcal{X})] f(\mathbf{x}_i) = \mathbf{y}^\mathsf{T}\mathbf{K}^{-1}\mu_k(\mathcal{X}),$$
$$e(Q_k)^2 = \int_\Omega \int_\Omega k(\mathbf{x},\mathbf{x}')\,\mathrm{d}\mu(\mathbf{x})\,\mathrm{d}\mu(\mathbf{x}') - \mu_k(\mathcal{X})^\mathsf{T}\mathbf{K}^{-1}\mu_k(\mathcal{X}) = \mu(\mu_k) - Q(\mu_k),$$

where $\mathbf{y} = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_n))$.

An *optimal kernel quadrature rule* minimises the worst-case error also over the node set (of fixed cardinality). Such rules cannot be constructed efficiently at the moment. We discuss structurally constrained versions in Section 4.3.

2.2. Probabilistic interpretation. The important probabilistic interpretation of kernel quadrature is a part of the field of emergent *probabilistic numerical computing* [12, 44, 23, 7], origins of which can be traced at least back to the work of Larkin [29]. This interpretation is a major motivator behind the present article.

In *Bayesian quadrature*, the integrand f is typically modelled as a Gaussian process [42, 49] (in which case the term *Gaussian process quadrature* is occasionally used) with the covariance kernel k. With the node locations $\mathbf{x}_1, \dots, \mathbf{x}_n$ and function evaluations $f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)$ considered the "data" $\mathcal{D}$, the posterior $f \mid \mathcal{D}$ is a Gaussian process with the mean and covariance

$$\mathbb{E}[f(\mathbf{x}) \mid \mathcal{D}] = \mathbf{y}^\mathsf{T}\mathbf{K}^{-1}k(\mathcal{X},\mathbf{x}),$$
$$\mathbb{C}[f(\mathbf{x}), f(\mathbf{x}') \mid \mathcal{D}] = k(\mathbf{x},\mathbf{x}') - k(\mathbf{x},\mathcal{X})\mathbf{K}^{-1}k(\mathcal{X},\mathbf{x}'),$$

where $[k(\mathcal{X},\mathbf{x})]_i = k(\mathbf{x}_i,\mathbf{x})$. This induces the Gaussian (because μ is a linear operator) posterior distribution $\mu(f) \mid \mathcal{D}$ on the integral with the mean $\mathbb{E}[\mu(f) \mid \mathcal{D}]$ and variance $\mathbb{V}[\mu(f) \mid \mathcal{D}]$ that turn out to be precisely $Q_k(f)$ and $e(Q_k)^2$ of the preceding section. The worst-case error can be therefore interpreted as a measure of uncertainty over the integral approximation and then exploited in for instance uncertainty quantification and allocation of limited computational resources in computational pipelines [7, Section 5]. Clear expositions of this Bayesian viewpoint to numerical integration are [43, 33, 4] and the methodology is quite popular in machine learning, see for example [48, 45].

3. Fully symmetric kernel quadrature. This is the main section of the article. We introduce fully symmetric sets, their connection to multivariate quadrature rules and prove our main result (Theorem 3.6) on computational benefits of doing kernel quadrature with node sets that are unions of fully symmetric sets.

3.1. Fully symmetric sets. A fully symmetric set is a point set in $\mathbb{R}^d$ that is obtained from a given vector through permutations and sign changes of its coordinates. Let Π_d be the set of all permutations $\mathbf{q} = (q_1, \dots, q_d)$ of the integers $1, \dots, d$ and S_d the set of all vectors of the form $\mathbf{s} = (s_1, \dots, s_d)$ with each s_i either 1 or -1. Then, given d non-negative scalars $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_d)$ called *generators*, the point set

$$[\boldsymbol{\lambda}] = [\lambda_1, \dots, \lambda_d] \coloneqq \bigcup_{\mathbf{q}\in\Pi_d} \bigcup_{\mathbf{s}\in S_d} \{(s_1\lambda_{q_1}, \dots, s_d\lambda_{q_d})\} \subset \mathbb{R}^d \tag{3.1}$$

is the *fully symmetric set* generated by the *generator vector* $\boldsymbol{\lambda}$. With m the number of non-zero generators and $m_1, \dots, m_l$ multiplicities of distinct generators, cardinality of the fully symmetric set (3.1) is

$$\#[\lambda_1, \dots, \lambda_d] = \frac{2^m d!}{m_1! \cdots m_l!}. \tag{3.2}$$

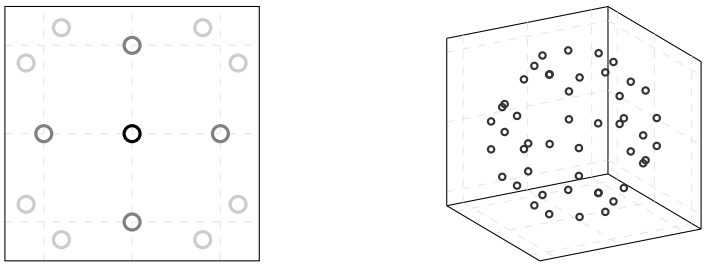

Fig. 3.1: Examples of fully symmetric sets in two and three dimensions. Left: the fully symmetric sets $[0,0]$, $[1,0]$ and $[1.2,0.8]$ in $\mathbb{R}^2$. Right: the fully symmetric set $[1,0.5,0.2]$ that consists of 48 elements in $\mathbb{R}^3$.

We use the convention that generators that are zero need not be specified. That is, $[\lambda_1,\dots,\lambda_p] = [\lambda_1,\dots,\lambda_p,0,\dots,0]$ when $p < d$.

An alternative way of writing Equation (3.1) is via *permutation matrices* as $[\boldsymbol{\lambda}] = \bigcup_{\mathbf{P}} \mathbf{P}\boldsymbol{\lambda}$ where the the union is over all $d \times d$ permutation and sign change matrices $\mathbf{P}$. These are matrices that have on each row and column exactly one element that is either 1 or -1 and the rest are zero. Any element of a fully symmetric set can be obtained from any other via linear transformation by an appropriate permutation matrix. How Equation (3.1) works and what the resulting point sets look like is illustrated in three and two dimensions in Example 3.2 and Figure 3.1. Note that all elements of a fully symmetric set are equidistant from the origin which is to say that if $\mathbf{x} \in [\lambda_1,\dots,\lambda_d]$, then $\|\mathbf{x}\|^2 = \|\boldsymbol{\lambda}\|^2 = \lambda_1^2 + \cdots + \lambda_d^2$. We also need the concept of a fully symmetric function.

DEFINITION 3.1. *A function $f\colon \Omega \to \mathbb{R}$ is* fully symmetric *if it is constant in every fully symmetric set. That is, with $\boldsymbol{\lambda}$ any generator vector, it holds that $f(\mathbf{x}) = f(\mathbf{x}')$ for any $\mathbf{x}, \mathbf{x}' \in \Omega \cap [\boldsymbol{\lambda}]$. Alternatively, $f(\mathbf{P}\mathbf{x}) = f(\mathbf{x})$ for any $\mathbf{x} \in \Omega$ and any permutation and sign change matrix $\mathbf{P}$ such that $\mathbf{P}\mathbf{x} \in \Omega$.*

EXAMPLE 3.2. *In $\mathbb{R}^3$, the non-zero and distinct generators λ_1 and λ_2 generate the fully symmetric set*

$$
[\lambda_1,\lambda_2,0] = \begin{array}{llll}
\{(\lambda_1,\lambda_2,0), & (-\lambda_1,\lambda_2,0), & (\lambda_1,-\lambda_2,0), & (-\lambda_1,-\lambda_2,0),\\
(\lambda_2,\lambda_1,0), & (-\lambda_2,\lambda_1,0), & (\lambda_2,-\lambda_1,0), & (-\lambda_2,-\lambda_1,0),\\
(0,\lambda_1,\lambda_2), & (0,-\lambda_1,\lambda_2), & (0,\lambda_1,-\lambda_2), & (0,-\lambda_1,-\lambda_2),\\
(0,\lambda_2,\lambda_1), & (0,-\lambda_2,\lambda_1), & (0,\lambda_2,-\lambda_1), & (0,-\lambda_2,-\lambda_1),\\
(\lambda_1,0,\lambda_2), & (-\lambda_1,0,\lambda_2), & (\lambda_1,0,-\lambda_2), & (-\lambda_1,0,-\lambda_2),\\
(\lambda_2,0,\lambda_1), & (-\lambda_2,0,\lambda_1), & (\lambda_2,0,-\lambda_1), & (-\lambda_2,0,-\lambda_1)\}
\end{array}
$$

that has $2^2 \times 3!/(1! \times 1! \times 1!) = 24$ elements. In terms of permutation matrices, the element $(-\lambda_1, 0, \lambda_2)$ is

$$
\begin{pmatrix} -\lambda_1 \\ 0 \\ \lambda_2 \end{pmatrix} = \begin{pmatrix} -1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ 0 \end{pmatrix}.
$$

3.2. Fully symmetric quadrature rules. The notation $f[\boldsymbol{\lambda}] = f[\lambda_1, \ldots, \lambda_d]$ stands for the sum of evaluations of f at the points of the fully symmetric set:

$$f[\boldsymbol{\lambda}] \coloneqq \sum_{\mathbf{x} \in [\boldsymbol{\lambda}]} f(\mathbf{x}).$$

A *fully symmetric quadrature rule* is a quadrature rule of the form

$$Q(f) = \sum_{\boldsymbol{\lambda} \in \Lambda} w_{\boldsymbol{\lambda}} f[\boldsymbol{\lambda}]$$

where Λ is a given finite collection of distinct generator vectors $\boldsymbol{\lambda}$. Such a rule uses only $\#\Lambda$ distinct weights, each corresponding to usually a very large number of nodes.

Fully symmetric quadrature rules are prominent among the classical polynomial quadrature rules, work on them going back to [31, 32]. To the best of our knowledge, the most general and efficient constructions have been given by Genz [16] for the uniform distribution on a square and Genz and Keister [17] for Gaussians on the whole real space (a case studied also in [30]). See for example the review [9] for more examples and discussion on the highly related invariant theory. To achieve high algebraic order of precision, the classical fully symmetric quadrature rules rely on symmetry of the underlying measure and advantageous properties of polynomials when integrated with respect to such measures. However, the rules become somewhat inflexible and the number of nodes needed often grows unnecessarily quickly.

Many of the popular sparse grid rules are also fully symmetric [40, 39]. We make use of this crucial fact in Section 4.2 for construction of sparse grid kernel quadrature rules whose weights can be computed efficiently.

3.3. Fully symmetric kernels. We can now introduce the class of kernels that this article is concerned with and the necessary assumptions on the integration domain and measure.

Definition 3.3. *A kernel k is* fully symmetric *if $k(\mathbf{P}\mathbf{x}, \mathbf{P}\mathbf{x}') = k(\mathbf{x}, \mathbf{x}')$ for any permutation and sign change matrix $\mathbf{P}$.*

This class of kernels includes (i) isotropic kernels, (ii) products of an isotropic kernel k_1 of the form $k(\mathbf{x}, \mathbf{x}') = \prod_{i=1}^d k_1(|x_i - x_i'|)$, and (iii) sums of an isotropic kernel k_1 of the form $k(\mathbf{x}, \mathbf{x}') = \sum_{i=1}^d k_1(|x_i - x_i'|)$. Polynomial kernels of the form $k(\mathbf{x}, \mathbf{x}') = \sum_{i=1}^p P_i(\mathbf{x}) P_i(\mathbf{x}')$ for suitable multivariate polynomials P_i are also fully symmetric. We believe that most classical interpolatory quadrature rules can be replicated if proper care is taken in choosing these polynomials.

Assumption 3.4. *We assume that*

1. *The integration domain $\Omega \subset \mathbb{R}^d$ is invariant under permutations and sign changes of coordinates of its elements. That is, $\Omega = \mathbf{P}\Omega = \{\mathbf{P}\omega : \omega \in \Omega\}$ for any permutation and sign change matrix $\mathbf{P}$.*
2. *The measure μ is fully symmetric in the sense that its density w (w.r.t. the Lebesgue measure) is a fully symmetric function.*
3. *The kernel k is fully symmetric.*

This assumption holds, for example, for $\Omega = [-1, 1]^d$ equipped with the uniform measure and $\Omega = \mathbb{R}^d$ equipped with the Gaussian measure as well as for many other cases of interest. The numerical examples in Section 5 are for these two cases.

Lemma 3.5. *The kernel mean is fully symmetric under Assumption 3.4.*

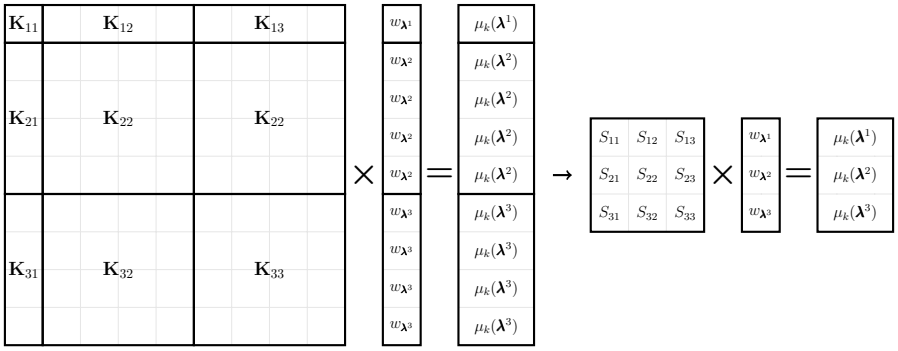
$\mathbf{K}_{11}$ $\mathbf{K}_{12}$ $\mathbf{K}_{13}$
$\mathbf{K}_{21}$ $\mathbf{K}_{22}$ $\mathbf{K}_{22}$
$\mathbf{K}_{31}$ $\mathbf{K}_{32}$ $\mathbf{K}_{33}$
$w_{\boldsymbol{\lambda}^1}$ $w_{\boldsymbol{\lambda}^2}$ $w_{\boldsymbol{\lambda}^2}$ $w_{\boldsymbol{\lambda}^2}$ $w_{\boldsymbol{\lambda}^2}$ $w_{\boldsymbol{\lambda}^3}$ $w_{\boldsymbol{\lambda}^3}$ $w_{\boldsymbol{\lambda}^3}$ $w_{\boldsymbol{\lambda}^3}$
$\mu_k(\boldsymbol{\lambda}^1)$ $\mu_k(\boldsymbol{\lambda}^2)$ $\mu_k(\boldsymbol{\lambda}^2)$ $\mu_k(\boldsymbol{\lambda}^2)$ $\mu_k(\boldsymbol{\lambda}^2)$ $\mu_k(\boldsymbol{\lambda}^3)$ $\mu_k(\boldsymbol{\lambda}^3)$ $\mu_k(\boldsymbol{\lambda}^3)$ $\mu_k(\boldsymbol{\lambda}^3)$
S_{11} S_{12} S_{13} S_{21} S_{22} S_{23} S_{31} S_{32} S_{33}
$w_{\boldsymbol{\lambda}^1}$ $w_{\boldsymbol{\lambda}^2}$ $w_{\boldsymbol{\lambda}^3}$
$\mu_k(\boldsymbol{\lambda}^1)$ $\mu_k(\boldsymbol{\lambda}^2)$ $\mu_k(\boldsymbol{\lambda}^3)$

Fig. 3.2: Illustration of Theorem 3.6 for a node set that is a union of three fully symmetric sets: one containing one element and two containing four elements. All row sums of $\mathbf{K}_{ij}$ are equal to S_{ij}.

Proof. Let $\mathbf{x}$ and $\mathbf{x}'$ be elements of the same fully symmetric set. That is, $\mathbf{x}' = \mathbf{P}\mathbf{x}$ for some permutation and sign change matrix $\mathbf{P}$. A change of variables yields

$$\int_\Omega k(\mathbf{x}', \mathbf{z}) w(\mathbf{z}) \, \mathrm{d}\mathbf{z} = \int_\Omega |\det \mathbf{P}| \, k\big(\mathbf{P}\mathbf{x}, \mathbf{P}\mathbf{z}\big) w\big(\mathbf{P}\mathbf{z}\big) \, \mathrm{d}\mathbf{z} = \int_\Omega k(\mathbf{x}, \mathbf{z}) w(\mathbf{z}) \, \mathrm{d}\mathbf{z},$$

where we have used the fact that $|\det \mathbf{P}| = 1$. That is, $\mu_k(\mathbf{x}') = \mu_k(\mathbf{x})$. ∎

3.4. Fully symmetric kernel quadrature. Let $[\boldsymbol{\lambda}^1], \ldots, [\boldsymbol{\lambda}^J]$ be distinct fully symmetric sets generated by $\boldsymbol{\lambda}^1, \ldots, \boldsymbol{\lambda}^J$. If the node set $\mathcal{X}$ is the union $\mathcal{X} = \cup_{j=1}^J [\boldsymbol{\lambda}^j]$ of these fully symmetric sets, then a kernel quadrature rule using this node set is a fully symmetric quadrature rule in the sense of Section 3.2. Furthermore, its J distinct weights can be computed extremely efficiently when compared to naively solving the linear system (2.2) of $\#\mathcal{X} = n$ equations. This is formalised in the following theorem and proved below. Figure 3.2 illustrates the simplified weight computation process in the case of a node set that is a union of three fully symmetric sets.

THEOREM 3.6. *Suppose that Ω, μ, and k satisfy Assumption 3.4. If the node set is a union of J disjoint fully symmetric sets $[\boldsymbol{\lambda}^1], \ldots, [\boldsymbol{\lambda}^J]$, then the kernel quadrature rule Q_k is fully symmetric:*

$$Q_k(f) = \sum_{j=1}^J w_{\boldsymbol{\lambda}^j} f[\boldsymbol{\lambda}^j].$$

Furthermore, the J weights $w_{\boldsymbol{\lambda}^1}, \ldots, w_{\boldsymbol{\lambda}^J}$ corresponding to the fully symmetric sets can be solved from the non-singular linear system of J equations

$$\begin{pmatrix} S_{11} & \cdots & S_{1J} \\ \vdots & \ddots & \vdots \\ S_{J1} & \cdots & S_{JJ} \end{pmatrix} \begin{pmatrix} w_{\boldsymbol{\lambda}^1} \\ \vdots \\ w_{\boldsymbol{\lambda}^J} \end{pmatrix} = \begin{pmatrix} \mu_k(\boldsymbol{\lambda}^1) \\ \vdots \\ \mu_k(\boldsymbol{\lambda}^J) \end{pmatrix}, \tag{3.4}$$

where S_{ij} are the row sums

$$S_{ij} = \sum_{\mathbf{x} \in [\boldsymbol{\lambda}^j]} k(\mathbf{x}^i, \mathbf{x}), \quad \text{for any} \quad \mathbf{x}^i \in [\boldsymbol{\lambda}^i], \tag{3.5}$$

of the submatrices $\mathbf{K}_{ij}$ in Equation (3.6).

Proof. Let $\mathcal{X}$ be ordered such that all elements of a single fully symmetric set appear consecutively and the fully symmetric sets themselves are in ascending order in terms of their index j.

By Lemma 3.5, the kernel mean is fully symmetric and, consequently, the vector $\mu_k(\mathcal{X})$ contains only J distinct elements $\mu_k(\boldsymbol{\lambda}^1), \ldots, \mu_k(\boldsymbol{\lambda}^J)$ that are grouped so that the first $\#[\boldsymbol{\lambda}^1]$ elements are equal to $\mu_k(\boldsymbol{\lambda}^1)$, the next $\#[\boldsymbol{\lambda}^2]$ to $\mu_k(\boldsymbol{\lambda}^2)$ and so on. For notational convenience, we denote $n^i = \#[\boldsymbol{\lambda}^i]$ and enumerate each generator as $[\boldsymbol{\lambda}^i] = \{\mathbf{x}_1^i, \ldots, \mathbf{x}_{n^i}^i\}$. Consider then the kernel matrix $\mathbf{K}$ that can be partitioned into J^2 submatrices $\mathbf{K}_{ij}$ of dimensions $n^i \times n^j$, each containing all the kernel evaluations $k(\mathbf{x}^i, \mathbf{x}^j)$ for $\mathbf{x}^i \in [\boldsymbol{\lambda}^i]$ and $\mathbf{x}^j \in [\boldsymbol{\lambda}^j]$. That is,

$$\mathbf{K} = \begin{pmatrix} \mathbf{K}_{11} & \cdots & \mathbf{K}_{1J} \\ \vdots & \ddots & \vdots \\ \mathbf{K}_{J1} & \cdots & \mathbf{K}_{JJ} \end{pmatrix}, \text{ where } \mathbf{K}_{ij} = \begin{pmatrix} k(\mathbf{x}_1^i, \mathbf{x}_1^j) & \cdots & k(\mathbf{x}_1^i, \mathbf{x}_{n^j}^j) \\ \vdots & \ddots & \vdots \\ k(\mathbf{x}_{n^i}^i, \mathbf{x}_1^j) & \cdots & k(\mathbf{x}_{n^i}^i, \mathbf{x}_{n^j}^j) \end{pmatrix}. \tag{3.6}$$

Any row of any submatrix $\mathbf{K}_{ij}$ can be obtained from any other of its rows by a permutation of elements of the row. To confirm this, consider $\mathbf{x}_p^i$ and $\mathbf{x}_l^j$ for any $p \leq n^i$ and $l \leq n^j$. For any $p' \neq p$, there is a permutation matrix $\mathbf{P}$ such that $\mathbf{x}_{p'}^i = \mathbf{P}\mathbf{x}_p^i$ because fully symmetric sets are closed under transformations by such matrices. Selecting l' such that $\mathbf{x}_{l'}^j = \mathbf{P}\mathbf{x}_l^j$ and using the fact that the kernel is fully symmetric yields

$$k(\mathbf{x}_{p'}^i, \mathbf{x}_{l'}^j) = k(\mathbf{P}\mathbf{x}_p^i, \mathbf{P}\mathbf{x}_l^j) = k(\mathbf{x}_p^i, \mathbf{x}_l^j),$$

which implies that the rows are permutations of each other. Consequently, the row sums $S_{ij} \coloneqq \sum_{\mathbf{x} \in [\boldsymbol{\lambda}^j]} k(\mathbf{x}^i, \mathbf{x})$ of $\mathbf{K}_{ij}$ do not depend on $\mathbf{x}^i \in [\boldsymbol{\lambda}^i]$.

Suppose now that all the weights for each fully symmetric set are equal. That is, $w_p = w_{\boldsymbol{\lambda}^1}$ for all $p = 1, \ldots, n^1$, $w_p = w_{\boldsymbol{\lambda}^2}$ for $p = n^1 + 1, \ldots, n^1 + n^2$ and so on. For any i such that the ith node is in $[\boldsymbol{\lambda}^p]$, the ith row of the system (2.2) is

$$\mu_k(\boldsymbol{\lambda}^p) = \sum_{j=1}^{J} w_{\boldsymbol{\lambda}^j} \sum_{l=1}^{n^j} [\mathbf{K}_{pj}]_{pl} = \sum_{j=1}^{J} w_{\boldsymbol{\lambda}^j} S_{pj},$$

because the row sums of each of the submatrix $\mathbf{K}_{pj}$ are equal to S_{pj}. This means that the J distinct weights can be solved from the linear system of equations

$$\underbrace{\begin{pmatrix} S_{11} & \cdots & S_{1J} \\ \vdots & \ddots & \vdots \\ S_{J1} & \cdots & S_{JJ} \end{pmatrix}}_{=\mathbf{S}} \begin{pmatrix} w_{\boldsymbol{\lambda}^1} \\ \vdots \\ w_{\boldsymbol{\lambda}^J} \end{pmatrix} = \begin{pmatrix} \mu_k(\boldsymbol{\lambda}^1) \\ \vdots \\ \mu_k(\boldsymbol{\lambda}^J) \end{pmatrix}.$$

We are left to show that the assumptions on the weights of each fully symmetric set being equal is true. The weights are unique because the kernel matrix is non-singular and we know that if the weights for each fully symmetric set are equal, they can be solved from the above system. That is, if the matrix $\mathbf{S}$ is non-singular, the weights possess the structure we have assumed. This matrix is singular if and only if there is a non-zero vector $\mathbf{b} \in \mathbb{R}^J$ such that $\mathbf{S}\mathbf{b} = \mathbf{0}$. But this would imply that $\mathbf{K}\mathbf{c} = \mathbf{0}$ with the non-zero vector $\mathbf{c} \in \mathbb{R}^n$ having the same structure as the full kernel mean vector

Table 4.1: Cardinalities, as computed from Equation (3.2), of fully symmetric sets generated by $m = 1, \ldots, 9$ distinct non-zero generators for dimensions $d = 2, \ldots, 9$.

	Dimension							
m	2	3	4	5	6	7	8	9
1	4	6	8	10	12	14	16	18
2	8	24	48	80	120	168	224	288
3		48	192	480	960	1,680	2,688	4,032
4			384	1,920	5,760	13,440	26,880	48,384
5				3,840	23,040	80,640	215,040	483,840
6					46,080	322,560	1,290,240	3,870,720
7						645,120	5,160,960	23,224,320
8							10,321,920	92,897,280
9								185,794,560

$\mu_k(\mathcal{X})$ but with b_j in the place of $\mu_k(\boldsymbol{\lambda}^j)$ for $j = 1, \ldots, J$. Because $\mathbf{K}$ is non-singular, this is impossible and hence $\mathbf{S}$ too is non-singular. This concludes the proof. □

The number of fully symmetric sets is rarely more than a few hundreds at most if the dimension is larger than three (see Table 4.1) so solving the linear system (3.4) is not a computational bottleneck. Instead, it is usually the Jn kernel evaluations that take most time in weight computations.

4. Selection of the fully symmetric sets. This section presents three different approaches for constructing the node set as a union of fully symmetric sets. Of these the sparse grids of Section 4.2 are the most promising alternative. We also discuss convergence properties of some of the kernel quadrature rules we construct in Section 4.4.

4.1. Random generators. Arguably, the simplest approach, both conceptually and algorithmically, is to draw a number of generator vectors randomly from the underlying distribution. However, unless additional constraints are enforced, all the generators will be distinct and non-zero, resulting in unrealistic numbers of integrand evaluations needed if $d \geq 6$, as seen from Table 4.1. One could heuristically set some generators to zero to reduce the number of nodes but it is not entirely clear how this should be done. In any case, for at least $d < 6$, the random generator approach seems realistic. We call this method the *fully symmetric kernel Monte Carlo* (FSKMC).

Theorem 4.1 provides theoretical convergence guarantees and Section 5.3 demonstrates that the FSKMC can also work in practice. Nevertheless, the method does not seem very promising as it comes across that a large number of random generator vectors, and thus an even larger number of nodes, is required to capture the underlying distribution.

This approach bears some similarity to stochastic radial and spherical integration rules developed in [18, 19]. These rules are less flexible due to the usual constraints of integrating low-degree polynomials exactly and more involved in their implementation.

4.2. Sparse grids. An iterated quadrature rule of degree m based on a regular Cartesian product grid requires m^d nodes—a number that quickly becomes impractically large. Sparse grids that originate in the work of Smolyak [52] are "sparsified" product sets widely used in numerical integration [37, 21, 39, 40]. See also the general survey by Bungartz and Griebel [5] and [24] for a wealth of financial applications. The

sparse grid construction that we present below is not the most general possible as we work in the fully symmetric framework. More general constructions are contained in some of the aforementioned references. In this section we assume that $\Omega = [-a, a]^d$ for a possibly infinite $a > 0$.

Let $X^1 = \{0\}$ and $X^i \subset X^{i+1} \subset [-a, a]$ for $i > 1$ be finite, nested and symmetric (i.e. if $x \in X^i$, then $-x \in X^i$) point sets. Then the *sparse grid* of *level* $q \geq 1$ is the set

$$H(q, d) \coloneqq \bigcup_{|\alpha|=d+q} \big(X^{\alpha_1} \times \cdots \times X^{\alpha_d}\big),$$

where $\alpha \in \mathbb{N}^d$ is a d-dimensional multi-index. Note that the largest X^i that is needed for a sparse grid of level q is X^{q+1}. As the basis sets X^i are nested and symmetric, it is easy to see that $H(q, d)$ is union of fully symmetric sets and can be explicitly written so:

$$H(q, d) = \bigcup_{\substack{|\alpha|=d+q \\ \alpha_i \geq \alpha_{i+1}}} \bigcup_{\substack{\lambda_i \in X^{\alpha_i} \\ \lambda_i \geq 0}} [\lambda_1, \ldots, \lambda_d].$$

That is, Theorem 3.6 applies to sparse grids. We call the resulting kernel quadrature rules the *sparse grid kernel quadratures*.

We are left with selection of the nested point sets X^i. In polynomial-based sparse grid quadrature rules these sets come coupled with univariate quadrature rules whose weights are used to construct the final sparse grid weights but we are under no such restrictions. An obvious idea for selecting X^i would be to sequentially minimise the worst-case error in one dimension—provided that the kernel is one for which this makes sense, for example any of the three examples of fully symmetric kernels given in Section 3.3. Discussion on different sequential kernel quadrature methods (usually known as *sequential Bayesian quadrature*) can be found in for example [8, 25, 22, 3].

However, owing to difficulties in setting the kernel length-scale, we do not employ this selection scheme. Instead, we use (i) the Clenshaw–Curtis point sets, rather standard in sparse grid literature, for $\Omega = [-1, 1]^d$ and the uniform measure and (ii) nested sets formed out of Gauss–Hermite nodes in the Gaussian case:

(i) For $i > 1$ and $m_i = 2^{i-1} + 1$, the nested Clenshaw–Curtis sets are

$$X^i = \{x_1^i, \ldots, x_{m_i}^i\} \quad \text{with} \quad x_j^i = -\cos\left(\frac{\pi(j-1)}{m_i - 1}\right) \in [-1, 1].$$

The points x_j^i are the roots and extrema of Chebyshev polynomials. The corresponding sparse grid kernel quadrature rule is called *Clenshaw–Curtis sparse grid kernel quadrature* (CCSGKQ). Numerical results for this kernel quadrature are given in Section 5.4 and convergence for sufficiently smooth functions is the topic of Theorem 4.2.

(ii) In the *Gauss–Hermite sparse grid kernel quadrature* (GHSGKQ) we use the classical Gauss–Hermite nodes that are the roots of the Hermite polynomials

$$\mathrm{H}_p(x) = (-1)^p \exp\big(x^2/2\big) \frac{\mathrm{d}^p}{\mathrm{d}x^p} \exp\big(-x^2/2\big).$$

Given a level q, we generate the $2q + 1$ symmetric roots of H_{2q+1} and for $i = 1, \ldots, q + 1$ select

$$X^i = \text{ the } 2i - 1 \text{ smallest roots by absolute value.}$$

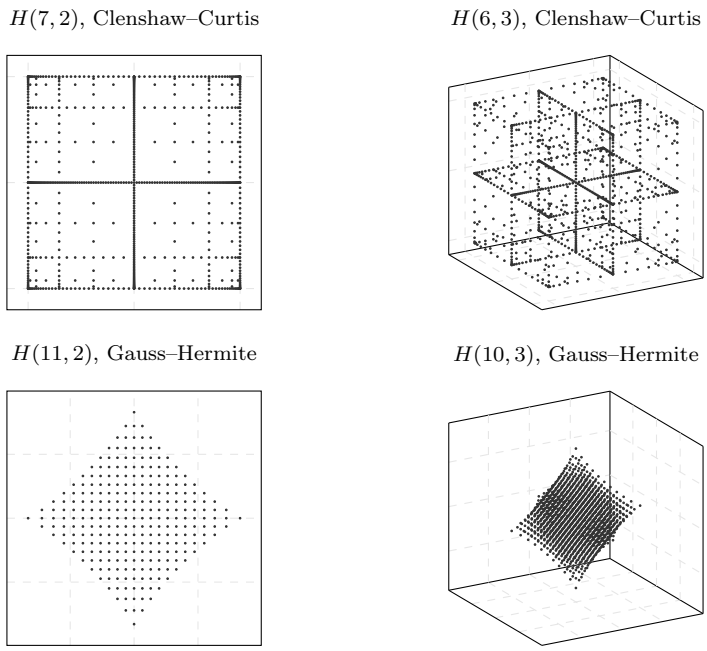

Fig. 4.1: Examples of sparse grids with Clenshaw–Curtis and Gauss–Hermite nodes. The numbers of nodes in the sparse grids are 705 (upper left), 1,073 (upper right), 265 (lower left), and 1,561 (lower right). Compare these to the cardinalities of the corresponding full grids that are $129^2 = 16{,}641$; $65^3 = 274{,}625$; $23^2 = 529$; and $21^3 = 9{,}261$.

The number of nodes, in terms of the level q, grows significantly slower than with Clenshaw–Curtis sparse grids. A numerical experiment involving a financial problem is given in Section 5.5. As is usual for quadrature rules on the whole of $\mathbb{R}^d$, there are no theoretical convergence guarantees for the GHSGKQ. Note that these grids are completely different from several other sparse grids in the literature that use nodes of Gaussian quadrature rules.

Four sparse grids based on these two point sequences are depicted in Figure 4.1. There is a large array of other possibilities available in the literature. For example, Gerstner and Griebel [21] use Gauss–Patterson nodes and Genz and Keister [17] have developed a nested version of the Gauss–Hermite rule. The rule (ii) can also be trivially extended for other integration domains and measures if a different sequence of orthogonal polynomials is used (e.g. Legendre or Chebyshev polynomials on $[-1, 1]$).

Finally, we remark that we are not the first to attempt combining kernel techniques and sparse grids—not even for numerical integration. See [15] for an application of sparse grids to machine learning and [20] for a kernel interpolation method based on them. The latter work has been applied, in a manner markedly different from ours, to numerical integration in [13, 55].

4.3. Worst-case error minimisation with respect to the generators. The third methodology for choosing the fully symmetric sets we briefly discuss (though do not test numerically) is that of principled RKHS worst-case error minimisation. Consider a fully symmetric kernel quadrature rule with the node set a union of $[\boldsymbol{\lambda}^1], \dots, [\boldsymbol{\lambda}^J]$. Suppose one, based on the number of nodes desired, fixes a number of

generators to zero or constraints them to be equal. Then the worst-case error $e(Q_k)$ of the kernel quadrature rule can be minimised with respect to the generator vectors obeying these constraints. Especially in higher dimensions, this is a task vastly simpler than trying to minimise the error over a node set of unconstrained geometry.

As a simple example, consider the fully symmetric rule

$$Q_k(f) = w_1 f[0] + w_2 f[\lambda_1, \lambda_2] = w_1 f[\mathbf{0}] + w_2 f[\lambda_1, \lambda_2] + w_3 f[\lambda_3, \lambda_4]$$

with $\lambda_1, \lambda_2, \lambda_3, \lambda_4 > 0$. If we apply the constraint $\lambda_3 = \lambda_4$, the optimal generators are

$$(\lambda_1^*, \lambda_2^*, \lambda_3^*, \lambda_4^*) = \underset{\substack{\lambda_1,\lambda_2,\lambda_3,\lambda_4>0 \\ \lambda_3=\lambda_4}}{\arg\min} e(Q_k).$$

4.4. Convergence analysis. In this section we provide convergence theorems for the fully symmetric kernel Monte Carlo and the Clenshaw–Curtis sparse grid kernel quadrature in terms of the worst-case error. The theorems are straightforward corollaries of some well-known results in the literature. For stating the results, we need to introduce the following three standard function classes. We assume that $\Omega = [-a, a]^d$ and $a < \infty$ in this section. The general principle on the convergence results for kernel quadrature is that the rates obtained are *at least as good as those for any other method using the same nodes if the integrand belongs to the RKHS induced by the kernel*. This should be quite clear from the definition of kernel quadrature.

With $\alpha \in \mathbb{N}^d$ a multi-index and $f \colon \Omega \to \mathbb{R}$ a sufficiently smooth function, the derivative operator is $\mathrm{D}^\alpha f = \partial^{|\alpha|} f / \partial x_1^{\alpha_1} \cdots \partial x_d^{\alpha_d}$. The function classes we need are

(i) The Sobolev space W_2^r is a Hilbert space defined as

$$W_2^r \coloneqq \left\{ f \in L^2(\mu) \,:\, \mathrm{D}^\alpha f \in L^2(\mu) \text{ exists for all } |\alpha| \leq r \right\}$$

with the norm $\|f\|_{W_2^r} = \sum_{|\alpha| \leq r} \|\mathrm{D}^\alpha f\|_{L^2(\mu)}$.

(ii) The class C^r is the class of functions that have bounded derivatives:

$$C^r \coloneqq \left\{ f \colon \Omega \to \mathbb{R} \,:\, \|\mathrm{D}^\alpha f\|_\infty < \infty \text{ for all } |\alpha| \leq r \right\}.$$

This space is equipped with the norm $\|f\|_{C^r} = \max\{\|\mathrm{D}^\alpha f\|_\infty \,:\, |\alpha| \leq r\}$.

(iii) The class F^r is the class of functions that have bounded mixed derivatives:

$$F^r \coloneqq \left\{ f \colon \Omega \to \mathbb{R} \,:\, \|\mathrm{D}^\alpha f\|_\infty < \infty \text{ for all } \alpha \leq r \right\}.$$

This space is equipped with the norm $\|f\|_{F^r} = \max\{\|\mathrm{D}^\alpha f\|_\infty \,:\, \alpha \leq r\}$.

For relations of the above function classes to reproducing kernel Hilbert spaces induced by different kernels, see [4, Section 3]. We merely note here a few connections that pertain slightly to the numerical examples of Section 5.

Two norms $\|\cdot\|$ and $\|\cdot\|'$ of an arbitrary vector space are *norm-equivalent* if there are positive constants C_1 and C_2 such that $C_1 \|x\| \leq \|x\|' \leq C_2 \|x\|$ for all elements x of the vector space. Recall from Section 2 that $\mathcal{H}$ is the RKHS induced by the kernel k. Then $\mathcal{H}$ is norm-equivalent to the Sobolev space W_2^r if $r > d/2$ and the Fourier transform $\mathcal{F}(\omega)$ of its kernel decays at the rate $(1 + \|\omega\|^2)^{-r}$. This holds if the kernel is for example of the Matérn class with $\nu \geq r - 1/2$ or if the kernel is Gaussian. In a similar manner, $\mathcal{H}$ is norm-equivalent to F^r if the kernel is a product of one-dimensional Matérn kernels.

The following convergence theorems are simple consequences of results available in the literature. Extensions are most likely possible using techniques in [26]. Concerning the rates of Theorem 4.2, see also the monographs [51, 41].

THEOREM 4.1 (Convergence of FSKMC). *Let $\Omega = [-a, a]^d$ with $a < \infty$. If $\mathcal{H}$ is norm-equivalent to the Sobolev space W_2^r with $r > d/2$, then*

$$\mathbb{E}\big[e(Q_k^{FSKMC})\big] = \mathcal{O}\big(n^{-r/d+\varepsilon}\big) \tag{4.1}$$

for any $\varepsilon > 0$. The expectation above is with respect to the joint distribution of the random generator vectors.

Proof. The worst-case error is decreasing in the number of nodes so we know that $e(Q_k^{\mathrm{FSKMC}}) \leq e(Q_k^{\mathrm{GEN}})$ where the rule Q_k^{GEN} uses only the J generator vectors as its nodes. The rate (4.1) is realised by the regular kernel Monte Carlo [4, Theorem 1] under the norm-equivalence assumption. Therefore,

$$\mathbb{E}\big[e(Q_k^{\mathrm{FSKMC}})\big] \leq \mathbb{E}\big[e(Q_k^{\mathrm{GEN}})\big] = \mathcal{O}\big(J^{-r/d+\varepsilon}\big) = \mathcal{O}\big(n^{-r/d+\varepsilon}\big)$$

because there is a dimension-dependent upper bound $2^d d!$ (see Equation (3.2)) for the number of nodes one fully symmetric set can contain. ∎

THEOREM 4.2 (Convergence of CCSGKQ). *Let $\Omega = [-a, a]^d$ with $a < \infty$ and μ be the uniform measure on Ω. If $\mathcal{H}$ is norm-equivalent to C^r, then*

$$e(Q_k^{CCSGKQ}) = \mathcal{O}\big(n^{-r/d}(\log n)^{(d-1)(r/d+1)}\big). \tag{4.2}$$

If $\mathcal{H}$ is norm-equivalent to F^r, then

$$e(Q_k^{CCSGKQ}) = \mathcal{O}\big(n^{-r}(\log n)^{(d-1)(r+1)}\big). \tag{4.3}$$

Proof. The rates (4.2) and (4.3) hold for the standard Clenshaw–Curtis sparse grid quadrature [37, 38] if the worst-case error (2.1) is over the unit balls of C^r and F^r, respectively. Because kernel quadrature rules have minimal worst-case errors in the induced RKHS among all quadrature rules with fixed nodes, the convergence rates follow from the assumptions of norm-equivalence. ∎

5. Numerical examples and computational aspects. This section contains three numerical examples for the fully symmetric kernel Monte Carlo, the Clenshaw–Curtis sparse grid kernel quadrature and the (modified) Gauss–Hermite sparse grid kernel quadrature as well as discussion on some computational aspects. The examples and algorithms, implemented in MATLAB, are available at `https://github.com/tskarvone/fskq`. Numerous classical sparse grid quadrature methods for MATLAB are implemented in the Sparse Grid Interpolation Toolbox [27]. Parts of our code make use of this toolbox.

We emphasise that none of the examples is meant to demonstrate superiority of fully symmetric kernel quadrature to other numerical integration methods. Comparisons to other methods are merely to show that fully symmetric kernel quadratures can achieve roughly comparable accuracy. Rather, we aim to show that fully symmetric sets make it possible to apply kernel quadrature rules to large-scale and high-dimensional situations that have been out of the scope of these quadrature rules before.

5.1. Choosing the length-scale parameter. Accuracy of any approximation based on a stationary kernel is heavily dependent on the length-scale parameter $\ell > 0$

whose effect is via $k_\ell(\mathbf{x}-\mathbf{x}') = k((\mathbf{x}-\mathbf{x}')/\ell)$, see for example the Gaussian kernel (5.1). Choosing in some sense the best value of this parameter efficiently is an important topic of research both in scattered data approximation literature [50, 14] and statistics and machine learning [49, Chapter 5]. See also [4, Section 4.1] for discussion in the context of kernel quadrature.

Unfortunately, we have been unable to apply any of the existing parameter fitting methods, such as marginal likelihood maximisation or cross-validation, to the present large-data scenario. Consequently, in the examples below we either use few enough nodes that naive computations are possible (Section 5.3), integrate a function whose length-scale is known beforehand (Section 5.4), or set the length-scale somewhat heuristically (Section 5.5). We recognise that the lack of a principled method for choosing the length-scale is a significant shortcoming and hope to fix this in the future.

5.2. Closed-form kernel means. In kernel quadrature, one needs to be able to evaluate the kernel mean $\mu_k(\mathbf{x}_i) = \int_\Omega k(\mathbf{x}_i, \mathbf{x}) \,\mathrm{d}\mu(\mathbf{x})$ at the nodes $\mathbf{x}_i$. A number of kernel-measure pairs that yield tractable kernel means are tabulated in [4]. It is also possible to evaluate the kernel mean numerically, see [4, Supplement C.4] and [53] for discussion and analysis of this case. Note that in fully symmetric setting numerical evaluation of the kernel mean would be quite viable as the kernel mean needs to be evaluated only at each generator vector instead of each node.

All our examples use the standard Gaussian kernel

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x}-\mathbf{x}'\|^2}{2\ell^2}\right) \tag{5.1}$$

with length-scale $\ell > 0$ and unit scale parameter (this parameter only affects magnitude of the worst-case error). The integration domain Ω and measure μ are either (i) the whole of $\mathbb{R}^d$ and the standard Gaussian measure with the density $\varphi(\mathbf{x}) = (2\pi)^{-d/2} \exp(-\|\mathbf{x}\|^2/2)$ or (ii) the hypercube $[-1,1]^d$ and the (normalising) uniform measure. In the former case the kernel mean and its integral (needed for the WCE) are

$$\mu_k(\mathbf{x}) = \int_{\mathbb{R}^d} k(\mathbf{x}, \mathbf{x}')\varphi(\mathbf{x}') \,\mathrm{d}\mathbf{x}' = \left(\frac{\ell^2}{1+\ell^2}\right)^{d/2} \exp\left(-\frac{\|\mathbf{x}\|^2}{2(1+\ell^2)}\right),$$
$$\mu(\mu_k) = \int_{\mathbb{R}^d} \mu_k(\mathbf{x})\varphi(\mathbf{x}) \,\mathrm{d}\mathbf{x} = \left(\frac{\ell^2}{2+\ell^2}\right)^{d/2}$$

and in the latter

$$\mu_k(\mathbf{x}) = 2^{-d} \int_{[-1,1]^d} k(\mathbf{x}, \mathbf{x}') \,\mathrm{d}\mathbf{x}' = \left(\frac{\pi\ell^2}{8}\right)^{d/2} \prod_{i=1}^{d} \left[\operatorname{erf}\left(\frac{x_i+1}{\ell\sqrt{2}}\right) - \operatorname{erf}\left(\frac{x_i-1}{\ell\sqrt{2}}\right)\right],$$
$$\mu(\mu_k) = 2^{-d} \int_{[-1,1]^d} \mu_k(\mathbf{x}) \,\mathrm{d}\mathbf{x} = \left(\frac{\pi\ell^2}{8}\right)^{d/2} \left(\sqrt{\frac{2\ell^2}{\pi}}\left(\mathrm{e}^{-2/\ell^2} - 1\right) + 2\operatorname{erf}\left(\sqrt{2}/\ell\right)\right)^d,$$

where $\operatorname{erf}(x) = \pi^{-1/2} \int_{-x}^{x} \exp(-t^2) \,\mathrm{d}t$ is the standard error function.

5.3. Example 1: Random generators. Our first example is just a proof of concept to demonstrate that the fully symmetric kernel Monte Carlo (FSKMC) from Section 4.1 indeed works. We try numerically integrating the non-radial function

$$f(\mathbf{x}) = \exp\left(\sin(5\|\mathbf{x}\|)^2 - (x_1^2 + 0.5x_2^2 + 2x_3^4)\right) \tag{5.2}$$

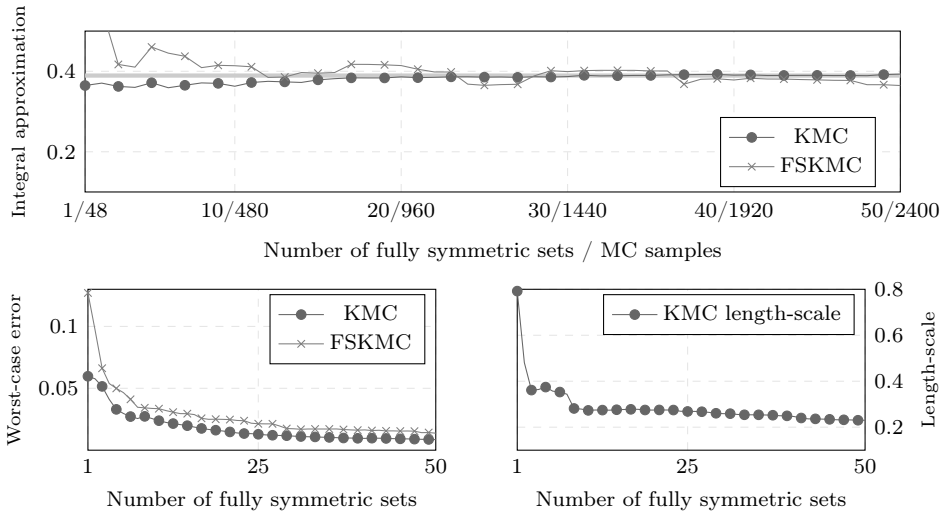

Fig. 5.1: Numerical integration of the function (5.2) with respect to the standard normal distribution. The upper figure shows integral approximations by the kernel Monte Carlo (KMC) and the fully symmetric kernel Monte Carlo (FSKMC) as a function of the number of fully symmetric sets or Monte Carlo samples for one realisation. Lower figures display the worst-case errors and the fitted length-scale. Each fully symmetric set contains 48 nodes (see Table 4.1). The underlying gray line is the true value of the integral which is approximately 0.389. The generator vectors for the FSKMC have been generated independently of the KMC samples. Both methods use the same length-scale that has been fit using the MC samples of the KMC.

over $\mathbb{R}^3$ and with respect to the standard normal distribution. Results for the kernel Monte Carlo (KMC) [48, 4], where the nodes to be used in kernel quadrature are drawn randomly, and the fully symmetric kernel Monte Carlo are presented in Figure 5.1.

For both the KMC and FSKMC, the kernel length-scale was fit by the method of maximum likelihood (see [49, Chapter 5]) using the MC samples of KMC. We have also experimented with fitting the FSKMC length-scale using the randomly generated fully symmetric sets, but in this case the fitted length-scale was markedly larger and the integral approximations much worse.

The results in show that in this example the FSKMC is able to produce similar result to the KMC, though its rate of convergence is slightly slower. However, the FSKMC has a tremendous advantage in computational scalability in the number of nodes. When the number of nodes is significantly over 10,000, kernel quadrature methods based on naively solving the weights from the linear system (2.2), such as the KMC, become infeasible due to their the cubic time and quadratic space complexity in the number of nodes. Fully symmetric kernel quadratures such as FSKMC are significantly more feasible: only Jn (recall that $J \leq n$ is the number of fully symmetric sets) kernel evaluations and solving a linear system of J equations, as opposed to n^2 and n, respectively, of naive methods, are required. For instance, using FSKMC with 1,000 fully symmetric sets (i.e. 48,000 nodes) in this example would require 48,000,000 kernel evaluations and solving a linear system of 1,000 equations, neither of which is a computational challenge, while the KMC weights for 48,000 nodes cannot be computed on a standard computer.

The difference becomes even more marked in higher dimensions where fully

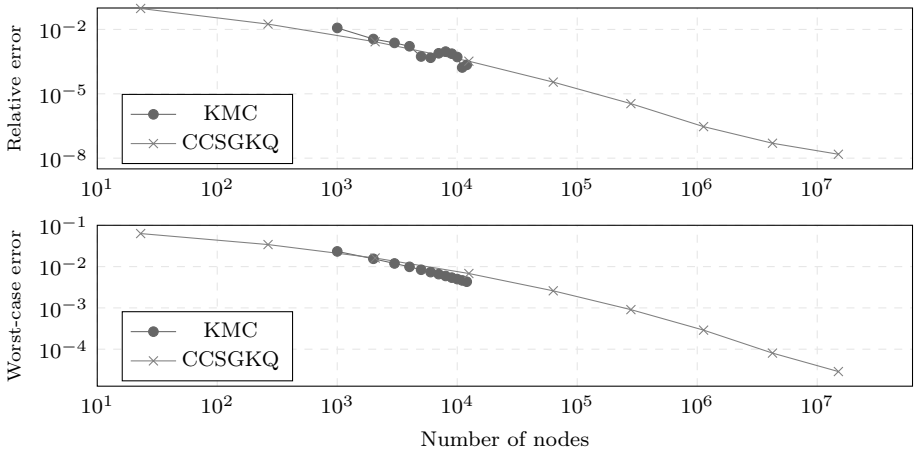

Fig. 5.2: Relative error $|\mu(f) - Q_k(f)| / \mu(f)$ (upper) and the worst-case error (lower) for integration of the function (5.3) on the 11-dimensional hypercube with the kernel Monte Carlo quadrature (KMC) and the Clenshaw–Curtis sparse grid kernel quadrature (CCSGKQ). The number of nodes for the KMC varied from 1,000 to 12,000 (with increments of 1,000) and CCSGKQ used levels from 1 to 9. These corresponded to 23; 265; 2,069; 12,497; 63,097; 280,017; 1,129,569; 4,236,673; and 15,005,761 nodes and 2, 4, 8, 17, 36, 79, 172, 379, and 832 fully symmetric sets.

symmetric sets contain more points. The next two examples demonstrate the superior scalability of fully symmetric kernel quadratures.

5.4. Example 2: *A priori* known length-scale. This simple example demonstrates that sparse grid kernel quadrature is numerically stable, consistent and works well for an extremely large number of nodes.

We work in the domain $\Omega = [-1, 1]^d$, $d = 11$, equipped with the normalising uniform measure. The integrand is

$$f(\mathbf{x}) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_f\|^2}{2\ell_f^2}\right) \tag{5.3}$$

with $\ell_f = 0.8$ and $\mathbf{x}_f$ is a vector of 11 evenly spaced points on the interval $[0.2, 0.5]$ (with the end points included). True value of this integral is

$$2^{-d} \int_{[-1,1]^d} f(\mathbf{x}) \, \mathrm{d}\mathbf{x} = \left(\frac{\pi \ell_f^2}{8}\right)^{d/2} \prod_{i=1}^{d} \left[\operatorname{erf}\left(\frac{x_{f,i} + 1}{\ell_f \sqrt{2}}\right) - \operatorname{erf}\left(\frac{x_{f,i} - 1}{\ell_f \sqrt{2}}\right)\right] \approx 0.0392.$$

We use the Gaussian kernel with $\ell = \ell_f = 0.8$ and experiment the Clenshaw–Curtis sparse grid kernel quadrature (CCSGKQ). Results for the *relative error* $|\mu(f) - Q_k(f)| / \mu(f)$ and the kernel RKHS worst-case error are shown in Figure 5.2 for the levels $q = 1, \ldots, 9$, last of them corresponding to the total of 15,005,761 nodes. We also display results for the kernel Monte Carlo with up to 12,000 nodes as after this solving the weights starts to become too time-intensive (see Section 5.3).

It is seen that for a similar numbers of nodes the two methods are roughly equivalent. When the level, and consequently the number of nodes, increases, the CCSGKQ becomes more accurate which shows that the nodes are selected well enough and that the weights are being computed correctly. For the highest levels 8 and 9, the sparse

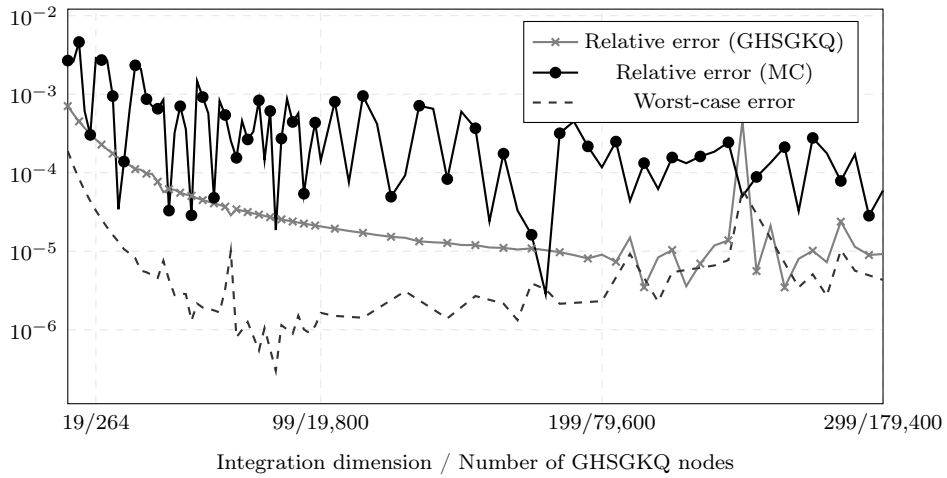

Fig. 5.3: Relative error and the worst-case error of the Gauss–Hermite sparse grid kernel quadrature (GHSGKQ) for the zero coupon bond setup of Section 5.5. True value of the integral (5.4), as computed from Equation (5.5), is between 0.81 and 0.815 for all dimensions. Also depicted is an integral estimate by the standard Monte Carlo using the same number of points as the GHSGKQ in each dimension.

grids consisted of 379 and 832 fully symmetric sets or 4,236,673 and 15,005,761 nodes, resulting in $379 \times 4{,}236{,}673 = 1{,}605{,}699{,}067$ and $832 \times 15{,}005{,}761 = 12{,}484{,}793{,}152$ kernel evaluations needed to compute the weights. On a fairly standard laptop it took about 2 and 17 minutes to compute the CCSGKQ weights for levels 8 and 9. It is not possible to compute the KMC weights for this many nodes.

5.5. Example 3: Zero coupon bonds. This example demonstrates that fully symmetric kernel quadrature rules are also feasible in high dimensions. We use the toy example from [36] (see also [24, Section 6.1]) that is concerned with pricing zero coupon bonds through simulation of a discretised stochastic differential equation model. The model is convenient for our purposes as there is a closed-form solution that serves as a baseline and finer discretisations correspond to higher integration dimensions.

Consider the stochastic differential equation (known as the Vasicek model)

$$\mathrm{d}r(t) = \kappa\big(\theta - r(t)\big)\,\mathrm{d}t + \sigma\,\mathrm{d}W(t),$$

where $W(t)$ is the standard Brownian motion and κ, θ, and σ are positive parameters. We want to solve this SDE at the time $t = T$. The Euler–Maruytama discretisation with the uniform step size $\Delta t = T/d$ is

$$r_k = r_{k-1} + \kappa(\theta - r_{k-1})\Delta t + \sigma x_k, \quad k = 1, \dots, d,$$

where $x_k \sim \mathcal{N}(0, \Delta t)$ are independent and r_0 is a free parameter. The quantity we are interested in is the Gaussian integral

$$\begin{aligned} P(0,T) &:= \mathbb{E}\bigg[\exp\bigg(-\Delta t \sum_{k=0}^{d-1} r_k\bigg)\bigg] \\ &= \exp(-\Delta t r_0) \int_{\mathbb{R}^{d-1}} \exp\Big(-\Delta t f\big(\sqrt{\Delta t}\mathbf{x}\big)\Big)\varphi(\mathbf{x})\,\mathrm{d}\mathbf{x}, \end{aligned} \tag{5.4}$$

where $f(\mathbf{x}) = \sum_{k=1}^{d-1} r_k$. This integral admits a closed-form solution

$$P(0,T) = \exp\left(-\frac{(\gamma + \beta_d r_0)T}{d} \right) \tag{5.5}$$

with $\beta_k = \sum_{i=1}^{k}(1-\kappa\Delta t)^{j-1}$ and $\gamma = \sum_{k=1}^{d-1}(\beta_k \kappa\theta\Delta t - (\beta_k\sigma\Delta t)^2/2)$. As can be seen, the number d of discretisation steps controls the integration dimension that is $d-1$.

In the integration experiment, we set

$$\kappa = 0.1817303, \quad \theta = 0.0825398957, \quad \sigma = 0.0125901, \quad r_0 = 0.021673, \quad T = 5.$$

These values are equal to those used in [36, 24]. We consider numerical integration of (5.4) for $d = 10, \ldots, 300$ using the Gauss–Hermite sparse grid kernel quadrature (GHSGKQ) with $q = 2$. We use the Gaussian kernel with the somewhat heuristic choice $\ell = d$. The central node (i.e. the origin) tended to have a fairly large negative weight so it was removed to improve numerical stability. Results for the relative error and the worst-case error are depicted in Figure 5.3. For comparison, we have also included a Monte Carlo estimate (KMC is feasible only for dimensions somewhat less than 100 so it was excluded, see again Section 5.3).

The results show that the GHSGKQ is able to maintain an accuracy that is generally better than that of the standard MC. This indicates that fully symmetric kernel quadratures have potential also in very high dimensions. Note the dimension-adaptive methods used in [24] would be more accurate in this example. It is probable that fully symmetric kernel quadratures could be combined with these methods.

6. Conclusions and discussion. We introduced fully symmetric kernel quadrature rules and showed that their weights can be computed exactly with a very simple algorithm under some assumptions on the integration domain and measure and the kernel. We also proposed using sparse grids in conjunction with this algorithm and provided some simple theoretical convergence analysis for Clenshaw–Curtis sparse grids. The nodes can be selected in a comparatively flexible manner. Three numerical experiments demonstrated that the approach is sound and can cope both with a very large number of nodes and a high integration dimension.

Even with the tremendous computational simplifications provided by the fully symmetric sets, kernel quadrature rules remain computationally more demanding than most classical quadrature rules. In the end, the decision on which method to use is highly dependent on the computational complexity of evaluating the integrand. Because the method presented in this article starts to become useful only after there are more than some thousands of nodes and may still incur large computational overhead compared to cheap integrand evaluations, we believe it is best suited for "moderately" expensive integrands in the case when high accuracy is required or probabilistic modelling of the error desired.

There are a number of topics that could be pursued in the future. These include

- Developing principled methods for choosing the kernel length-scale for large-scale problems.
- Proper probabilistic approach to large-scale integration problems. We anticipate that much can be gained in pursuing this direction.
- As discussed in Sections 4.2 and 4.3, there is much room for improvement via optimisation of the fully symmetric sets.
- Rows of the submatrices $\mathbf{K}_{ij}$ in Equation (3.6) typically contain several non-distinct elements. Minor computational improvements might be possible.

Acknowledgements. We thank Filip Tronarp for helpful comments.

REFERENCES

[1] N. Aronszajn, *Theory of reproducing kernels*, Transactions of the American Mathematical Society, 68 (1950), pp. 337–404.
[2] A. Y. Bezhaev, *Cubature formulae on scattered meshes*, Russian Journal of Numerical Analysis and Mathematical Modelling, 6 (1991), pp. 95–106.
[3] F.-X. Briol, C. J. Oates, M. Girolami, and M. A. Osborne, *Frank-Wolfe Bayesian quadrature: Probabilistic integration with theoretical guarantees*, in Advances in Neural Information Processing Systems, 2015, pp. 1162–1170.
[4] F.-X. Briol, C. J. Oates, M. Girolami, M. A. Osborne, and D. Sejdinovic, *Probabilistic integration: A role for statisticians in numerical analysis?*, (2016). Preprint, arXiv:1512.00933.
[5] H.-J. Bungartz and M. Griebel, *Sparse grids*, Acta Numerica, 13 (2004), pp. 147–269.
[6] R. E. Caflisch, *Monte Carlo and quasi-Monte Carlo methods*, Acta Numerica, 7 (1998), pp. 1–49.
[7] J. Cockayne, C. Oates, T. Sullivan, and M. Girolami, *Bayesian probabilistic numerical methods*, (2017). Preprint, arXiv:1702.03673.
[8] T. D. Cook and M. K. Clayton, *Sequential Bayesian quadrature*, tech. report, Department of Statistics, University of Wisconsin, Jul 1998.
[9] R. Cools, *Constructing cubature formulae: The science behind the art*, Acta Numerica, 6 (1997), pp. 1–54.
[10] P. J. Davis and P. Rabinowitz, *Methods of Numerical Integration*, Academic Press, 2nd ed., 1984.
[11] C. de Boor and A. Ron, *On multivariate polynomial interpolation*, Constructive Approximation, 6 (1990), pp. 287–302.
[12] P. Diaconis, *Bayesian numerical analysis*, in Statistical decision theory and related topics IV, vol. 1, Springer-Verlag New York, 1988, pp. 163–175.
[13] Z. Dong, E. H. Georgoulis, J. Levesley, and F. Usta, *Fast multilevel sparse Gaussian kernels for high-dimensional approximation and integration*, (2015). Preprint, arXiv:1501.03296.
[14] G. E. Fasshauer and J. G. Zhang, *On choosing "optimal" shape parameters for RBF approximation*, Numerical Algorithms, 45 (2007), pp. 345–368.
[15] J. Garcke, *A dimension adaptive sparse grid combination technique for machine learning*, ANZIAM Journal, 48 (2006), pp. 725–740.
[16] A. Genz, *Fully symmetric interpolatory rules for multiple integrals*, SIAM Journal on Numerical Analysis, 23 (1986), pp. 1273–1283.
[17] A. Genz and B. D. Keister, *Fully symmetric interpolatory rules for multiple integrals over infinite regions with Gaussian weight*, Journal of Computational and Applied Mathematics, 71 (1996), pp. 299–309.
[18] A. Genz and J. Monahan, *Stochastic integration rules for infinite regions*, SIAM Journal on Scientific Computing, 19 (1998), pp. 426–439.
[19] A. Genz and J. Monahan, *A stochastic algorithm for high-dimensional integrals over unbounded regions with Gaussian weight*, Journal of Computational and Applied Mathematics, 112 (1999), pp. 71–81.
[20] E. H. Georgoulis, J. Levesley, and F. Subhan, *Multilevel sparse kernel-based interpolation*, SIAM Journal on Scientific Computing, 35 (2013), pp. A815–A831.
[21] T. Gerstner and M. Griebel, *Numerical integration using sparse grids*, Numerical Algorithms, 18 (1998), pp. 209–232.
[22] T. Gunter, M. A. Osborne, R. Garnett, P. Hennig, and S. J. Roberts, *Sampling for inference in probabilistic models with fast Bayesian quadrature*, in Advances in Neural Information Processing Systems, 2014, pp. 2789–2797.
[23] P. Hennig, M. A. Osborne, and M. Girolami, *Probabilistic numerics and uncertainty in computations*, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, 471 (2015).
[24] M. Holtz, *Sparse Grid Quadrature in High Dimensions with Applications in Finance and Insurance*, no. 77 in Lecture Notes in Computational Science and Engineering, Springer, 2011.
[25] F. Huszár and D. Duvenaud, *Optimally-weighted herding is Bayesian quadrature*, in Proceedings of the Conference on Uncertainty in Artificial Intelligence, 2012, pp. 377–385.
[26] M. Kanagawa, B. K. Sriperumbudur, and K. Fukumizu, *Convergence guarantees for*

kernel-based quadrature rules in misspecified settings, in Advances in Neural Information Processing Systems, vol. 29, 2016, pp. 3288–3296.

[27] A. Klimke and B. Wohlmuth, *Algorithm 847: spinterp: Piecewise multilinear hierarchical sparse grid interpolation in MATLAB*, ACM Transactions on Mathematical Software, 31 (2005), pp. 561–579.

[28] F. M. Larkin, *Optimal approximation in Hilbert spaces with reproducing kernel functions*, Mathematics of Computation, 24 (1970), pp. 911–921.

[29] F. M. Larkin, *Gaussian measure in Hilbert space and applications in numerical analysis*, Rocky Mountain Journal of Mathematics, 2 (1972), pp. 379–422.

[30] J. Lu and D. L. Darmofal, *Higher-dimensional integration with Gaussian weight for applications in probabilistic design*, SIAM Journal on Scientific Computing, 26 (2004), pp. 613–624.

[31] J. N. Lyness, *Symmetric integration rules for hypercubes I. Error coefficients*, Mathematics of Computation, 19 (1965), pp. 260–276.

[32] J. McNamee and F. Stenger, *Construction of fully symmetric numerical integration formulas*, Numerische Mathematik, 10 (1967), pp. 327–344.

[33] T. Minka, *Deriving quadrature rules from Gaussian processes*, tech. report, Statistics Department, Carnegie Mellon University, Nov 2000.

[34] K. Muandet, K. Fukumizu, B. Sriperumbudur, and B. Schölkopf, *Kernel mean embedding of distributions: A review and beyond*, (2016). Preprint, arXiv:1605.09522.

[35] A. Narayan and D. Xiu, *Stochastic collocation methods on unstructured grids in high dimensions via interpolation*, SIAM Journal on Scientific Computing, 34 (2012), pp. A1729–A1752.

[36] S. Ninomiya and S. Tezuka, *Toward real-time pricing of complex financial derivatives*, Applied Mathematical Finance, 3 (1996), pp. 1–20.

[37] E. Novak and K. Ritter, *High dimensional integration of smooth functions over cubes*, Numerische Mathematik, 75 (1996), pp. 79–97.

[38] E. Novak and K. Ritter, *The curse of dimension and a universal method for numerical integration*, in Multivariate Approximation and Splines, Birkhäuser Basel, Nov 1997, pp. 177–187.

[39] E. Novak and K. Ritter, *Simple cubature formulas with high polynomial exactness*, Constructive approximation, 15 (1999), pp. 499–522.

[40] E. Novak, K. Ritter, R. Schmitt, and A. Steinbauer, *On an interpolatory method for high dimensional integration*, Journal of Computational and Applied Mathematics, 112 (1999), pp. 215–228.

[41] E. Novak and H. Woźniakowski, *Tractability of Multivariate Problems, volume II: Standard Information for Functionals*, no. 12 in EMS Tracts in Mathematics, European Mathematical Society, 2010.

[42] A. O'Hagan, *Curve fitting and optimal design for prediction*, Journal of the Royal Statistical Society. Series B (Methodological), 40 (1978), pp. 1–42.

[43] A. O'Hagan, *Bayes–Hermite quadrature*, Journal of Statistical Planning and Inference, 29 (1991), pp. 245–260.

[44] A. O'Hagan, *Some Bayesian numerical analysis*, Bayesian Statistics, 4 (1992), pp. 345–363.

[45] M. A. Osborne, D. Duvenaud, R. Garnett, C. E. Rasmussen, S. J. Roberts, and Z. Ghahramani, *Active learning of model evidence using Bayesian quadrature*, in Advances in Neural Information Processing Systems, 2012, pp. 46–54.

[46] J. Prüher and O. Straka, *Gaussian process quadrature moment transform*, (2017). Preprint, arXiv:1701.01356.

[47] J. Prüher, F. Tronarp, T. Karvonen, S. Särkkä, and O. Straka, *Student-t process quadratures for filtering of non-linear systems with heavy-tailed noise*, (2017). Preprint, arXiv:1703.05189.

[48] C. E. Rasmussen and Z. Ghahramani, *Bayesian Monte Carlo*, in Advances in Neural Information Processing Systems, 2002, pp. 489–496.

[49] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, Adaptive Computation and Machine Learning, MIT Press, 2006.

[50] S. Rippa, *An algorithm for selecting a good value for the parameter c in radial basis function interpolation*, Advances in Computational Mathematics, 11 (1999), pp. 193–210.

[51] K. Ritter, *Average-Case Analysis of Numerical Problems*, no. 1733 in Lecture Notes in Mathematics, Springer, 2000.

[52] S. A. Smolyak, *Quadrature and interpolation formulas for tensor products of certain classes of functions*, Doklady Akademii Nauk SSSR, 4 (1963), pp. 240–243.

[53] A. Sommariva and M. Vianello, *Numerical cubature on scattered data by radial basis functions*, Computing, 76 (2006), pp. 295–310.

[54] S. Särkkä, J. Hartikainen, L. Svensson, and F. Sandblom, *On the relation between Gaussian process quadratures and sigma-point methods*, Journal of Advances in Information Fusion, 11 (2016), pp. 31–46.

[55] F. Usta and J. Levesley, *Quasi-interpolation on a sparse grid with Gaussian*, (2016). Preprint, arXiv:1608.00728.